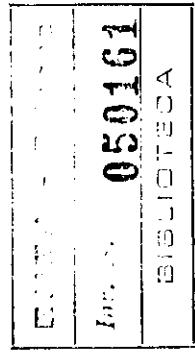


Fundamentals of Interactive Computer Graphics

1) Amplificazioni diverse lungo l'asse centrale
(taratura sul binario è importante!)

JAMES D. FOLEY
The George Washington University
ANDRIES VAN DAM
Brown University



Addison-Wesley Publishing Company
Reading, Massachusetts • Menlo Park, California
London • Amsterdam • Don Mills, Ontario • Sydney



Fig. 17.13 Photo reproduced with D^8 ordered dither on 2×2 display (courtesy J. Jarvis, Bell Laboratories).

17.3 CHROMATIC COLOR

The visual sensations caused by color are much richer than the sensations caused by achromatic light. Subjective discussions of color usually involve three dimensions, known as *hue*, *saturation*, and *brightness*, as a descriptive tool. *Hue* is the term we use to distinguish between colors such as red, green, yellow, etc. *Saturation* refers to purity, i.e., how little the color is diluted by white light, and distinguishes pink from red, sky blue from royal blue, etc. In other words, saturation determines how pastel or strong a color appears. Brightness embodies the achromatic notion of intensity, used in the preceding sections, as a factor independent of hue and saturation.

We are interested in specifying and measuring color. One way to do this is by visually comparing a sample of unknown color against a set of "standard" samples. The widely used Munsell color specification system includes sets of published standard colors [MUNS76] organized in a three-dimensional space of hue, value (brightness), and chroma (saturation). Each color has a name, and is an equal perceived "distance" in color space (as judged by many observers) from its neighbors. In [KELL76] there is an extensive discussion of standard samples, charts depicting the Munsell space, and tables of color names. The Ostwald [OSTW31] system is similar but somewhat less used. There is also the relatively new Coloroid system of Nemsics [NEMC80].

Artists use another approach, specifying color as different *tints*, *shades*, and tones of strongly saturated, or pure, pigments. A tint results from adding white pigment to the pure pigment, thereby decreasing saturation. A shade comes from adding a black pigment to the pure pigment, resulting in decreased brightness. A tone is the consequence of adding both black and white pigments to a pure pigment.

All these steps produce different colors of the same hue, with varying saturation and brightness. Mixing just black and white pigments creates grays. Figure 17.14 shows the relationships of tints, shades, and tones. We could think of using the percentage of pigments which must be mixed to match a color as a measurement of a color.

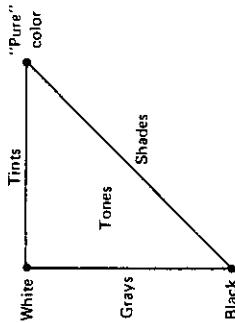


Fig. 17.14 Tints, tones, and shades.

17.3.1 Color in Physics and Physiology

Unfortunately, the Munsell and pigment mixing methods are subjective, i.e., depend on the observers' judgments as to color matches. What we really need is an objective way of specifying colors. For this we turn to physics, wherein visible light is treated as electromagnetic energy with a spectral energy distribution in the visible part of the spectrum, which ranges from violet through indigo, blue, green, yellow and orange, to red (remembered by the mnemonic VIB-GY-OR or, in reverse, Mr. ROY G. BIV). Figure 17.15 shows a typical spectral energy distribution of a light source, and represents an infinity of numbers, one for each wavelength in the visible spectrum (in reality, the distribution would be defined by a large number of sample points on the spectrum). Fortunately, we can represent the visual effect of any spectral distribution in a much more concise way, by the triple (dominant wavelength, purity, luminance). This implies that many different spectral energy distributions produce the same color; they "look" the same. This means that the relationship between spectral distributions and colors is many-to-one.

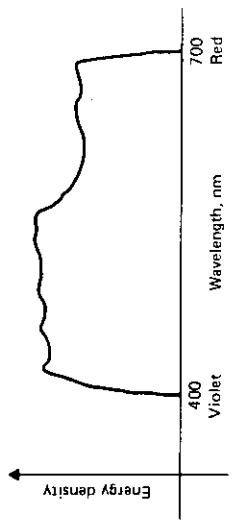


Fig. 17.15 Typical spectral energy distribution

The components of the triple have the following interpretations: *dominant wavelength* is the wavelength of the color we “see” when viewing the light, and corresponds to the subjective notion of hue; *purity* corresponds to saturation of the color; and *luminance* is the amount of light. For achromatic light, luminance is the light’s intensity. The purity of a colored light is the proportion of pure light of the dominant wavelength and white light needed to define the color. A completely pure color, which is thus 100% saturated, contains no white light. A half-and-half mixture would be 50% saturated. White light and hence all gray levels are 0% saturated, containing no color of any dominant wavelength.

Figure 17.16 shows one of the infinitely many spectral distributions which will produce a certain color of light. At the dominant wavelength there is a spike of energy of level e_2 . White light, represented by the uniform distribution of energy at level e_1 , is also present. Purity depends on the relation between e_1 and e_2 . When $e_1 = e_2$, purity is 0%. When $e_1 = 0$, purity is 100%. Luminance, which can be thought of as the area under the curve (total energy), depends on both e_1 and e_2 .

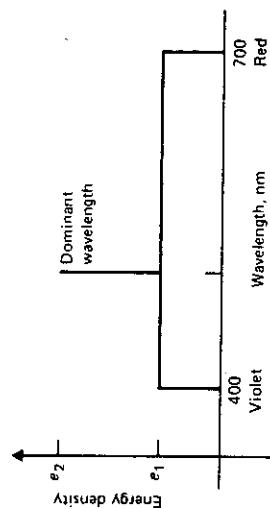


Fig. 17.16 Spectral energy distribution illustrating dominant wavelength, purity, and luminance.

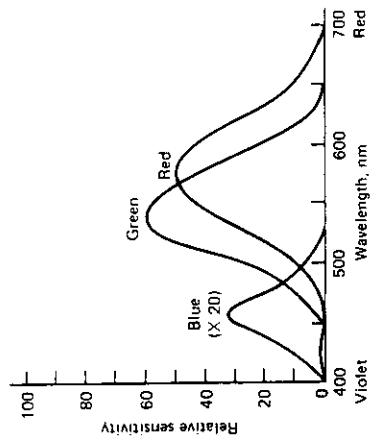


Fig. 17.17 Response characteristics of the eye.

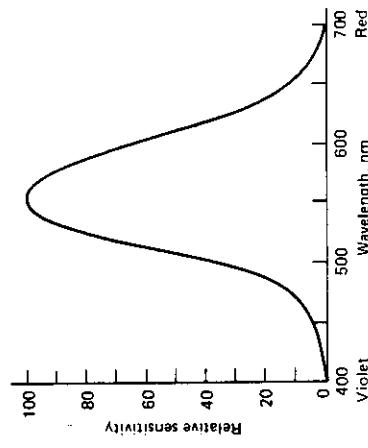


Fig. 17.18 Luminosity response of the eye.

Here then is a precise definition of color in terms of dominant wavelength, purity, and luminance. But how does this relate to the red, green, and blue phosphor dots on a color CRT and to the psychological-physiological *tri-stimulus theory* of color based on the hypothesis that there are three kinds of cones on the retina of the eye, each with peak sensitivities to light of either red, green, or blue hues? Experiments based on this hypothesis produce the response curves of Fig. 17.17. The curves show, for instance, that for 550 nanometer (nm) (also called millimicron) wavelength light, the blue receptors have a sensitivity of 0%; the green, about 55%; the red, about 45%. The curves also indicate that the blue receptors are far less sensitive than the red and green receptors. The sum of the three response curves, shown in Fig. 17.18, is known as the *luminosity curve*. It shows the eye’s response to light of constant luminance as the dominant wavelength is varied: our peak sensitivity is to yellow-green light of wavelength around 550 nm.

The tri-stimulus approach is not the only theory used to explain color vision. Land’s recent theories [LAND77], *opponent-color theory*, and *zone theory* are also used [WASS78]. The tri-stimulus theory, however, is attractive on an intuitive basis because it loosely corresponds to the notion that colors can be specified by weighted sums of red, green, and blue. This notion is in fact almost true: the three curves in Fig. 17.19 show the amounts of red (650 nm wavelength), green (530 nm), and blue (425 nm) light needed by an “average” observer to match a 100% pure light (a spectral color) of constant luminance, for all values of dominant wavelength in the visible spectrum.

Some colors actually cannot be matched, but it happens that by adding a primary to the color, it can then be matched by the other two primaries. A negative value in Fig. 17.19 indicates that the primary was added to the color being matched.

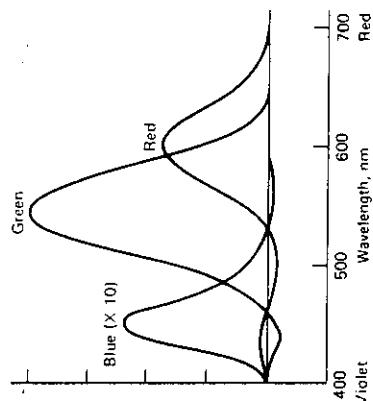


Fig. 17.19 Color matching coefficients. The blue primary has been scaled up by 10.

This does not mean that the notion of mixing red, green, and blue to obtain other colors is invalid; on the contrary, there is a huge gamut (range) of colors which can be matched by positive amounts of red, green, and blue. Otherwise color TV wouldn't work! It is important to notice that because the human eye is less sensitive to blue light than to green light, less blue light is required in the matching process.

The human eye can distinguish about 350,000 different colors. This figure is based on experiments in which many pairs of colors are judged side by side by many viewers. The viewer states whether the colors are the same or different. When the colors differ only in hue, then the wavelength between just noticeably different colors varies from greater than 10 nm at the extremes of the spectrum to about 1 nm for blue and yellow. Except at the spectrum extremes, most distinguishable colors are within 3 nm. Altogether about 128 hues are distinguishable. If colors differ only in saturation, we can distinguish from 16 (for yellow) to 23 (for red and violet).

17.3.2 The CIE Chromaticity Diagram

Matching and therefore defining a colored light with a mixture of three fixed primaries is a desirable approach, but the concept of negative weights suggested by Fig. 17.19 is unattractive. In 1931, the *Commission Internationale L'Eclairage* (CIE) defined three primary colors (X , Y , Z) that can be combined, with positive weights, to define all light sensations we experience with our eyes. The CIE primaries, which are by themselves not visible, form an international standard for specifying color. The primaries are defined as three spectral energy distributions: the Y primary was intentionally defined to have an energy distribution which exactly matches the luminosity curve of Fig. 17.18.

Let (X, Y, Z) be the weights applied to the CIE primaries to match a color. We can define chromaticity values (which depend only on dominant wavelength and

saturation and are independent of the amount of luminous energy) by normalizing against luminance (which is the total amount of light) as follows:

$$x = \frac{X}{X + Y + Z}, \quad y = \frac{Y}{X + Y + Z}, \quad z = \frac{Z}{X + Y + Z}. \quad (17.12)$$

Notice that we have forced $x + y + z$ to equal 1.

By plotting x and y for all visible colors, we obtain the CIE chromaticity diagram shown in Fig. 17.20. The interior and boundary of the horseshoe-shaped region represent all visible chromaticities. (All perceptible colors with the same chromaticity but different luminances map into the same point within this region.) The 100% pure colors of the spectrum are on the curved part of the boundary. Their wavelengths are indicated in the figure. Standard white light, meant to approximate sunlight, is formally defined by a standard light source known as *illuminant C* which is marked by the center dot. It is near but not at the point where $x = y = z = \frac{1}{3}$. Illuminant C is precisely defined in colorimetry (the science of color measurement) as a black-body radiator at 6504° Kelvin.

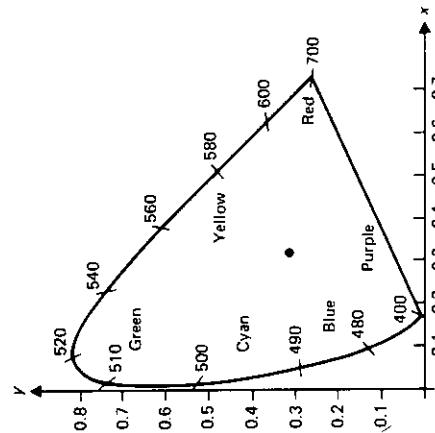


Fig. 17.20 CIE chromaticity diagram. Wavelengths are in nanometers.

The CIE chromaticity diagram is useful in many ways. For one, it allows us to actually measure the dominant wavelength and purity of any color, by first matching the color by a mixture of the three CIE primaries. (There are instruments which can do this). Now suppose the matched color is at point A in Fig. 17.21. When two colors are added together, the new color lies somewhere on the straight line in the chromaticity diagram connecting the two colors being added. Therefore color A can be thought of as a mixture of "standard" white light (illuminant C) and the pure spectral light at point B. Thus B defines the dominant wavelength. The ratio of

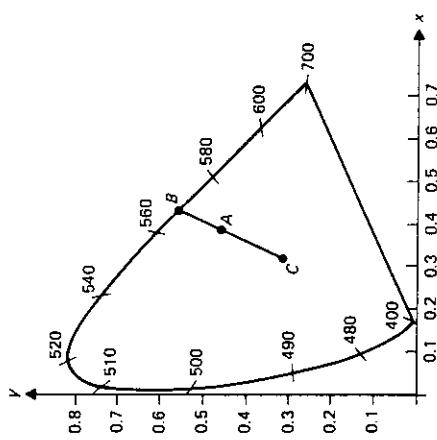


Fig. 17.21 Dominant wavelengths of color *A* is that of color *B* (about 565 nanometers).

length AC to length BC , expressed as a percentage, is the purity of *A*. The closer *A* is to *C*, the more white light *A* includes and the less pure it is. Because the diagram factors out luminance, color sensations which are luminance-related are excluded. Therefore brown, which is an orange-red chromaticity at very low luminance, is not shown.

Complementary colors are those that can be mixed to produce white light (such as *D* and *E* in Fig. 17.22). Here the weight w_D applied to chromaticity *D* in the mixture is the ratio of line length \overline{CE} to line length \overline{DE} , and the weight w_E applied to chromaticity *E* is the ratio of line length \overline{DC} to \overline{DE} ; hence $C = w_D D + w_E E$. Representing a chromaticity by its (x, y) coordinates, we have

$$x_C = w_D x_D + w_E x_E \quad \text{and} \quad y_C = w_D y_D + w_E y_E.$$

Notice that $w_D = 1 - w_E$; this is just the parametric representation of a line used in Chapters 4 and 8, with $t = w_D$.

Some colors (such as *E* in Fig. 17.23) cannot be defined by a dominant wavelength and are thus called *non-spectral*. In these cases, the dominant wavelength is said to be the complement of that at point *D* and is designated by a "c" (in this case, about 560 nm c). The purity is still defined from the ratio of distances (in this case, CE and CF). The colors that must be expressed by using a complementary dominant wavelength are the purples and the magentas; they occur in the lower part of the CIE diagram.

Another use for the CIE chromaticity diagram is to define *color gamuts*, or color ranges. Any two colors (*I* and *J* in Fig. 17.24) can be added to produce any color on the connecting line by varying the relative luminances of the two colors be-

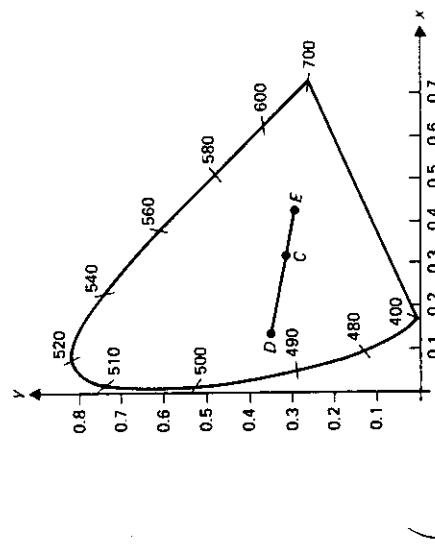


Fig. 17.22 Complementary colors.

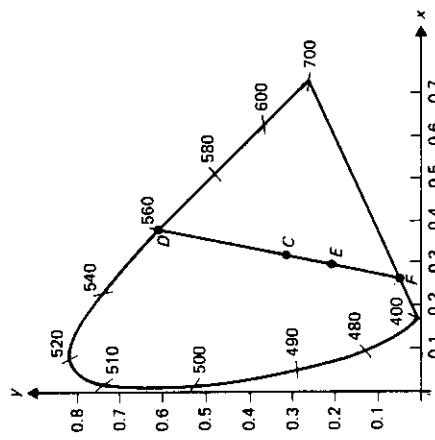


Fig. 17.23 The dominant wavelength of color *E* is defined as the complement of the dominant wavelength of color *D*.

ing added. A third color *K* (see Fig. 17.25) can be used with various mixtures of *I* and *J* to produce the gamut of all colors in triangle *IJK*, again by varying relative luminances. The shape of the diagram shows why visible red, green, and blue cannot be additively mixed to match all colors; no triangle whose vertices are within the visible area can completely cover the visible area. Any fully saturated color, because it contains no white light, can be specified as a mixture of just two primaries.

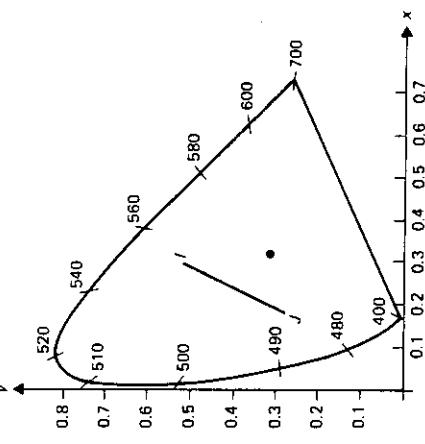


Fig. 17.24 Mixing two colors. All colors on connecting line can be created.

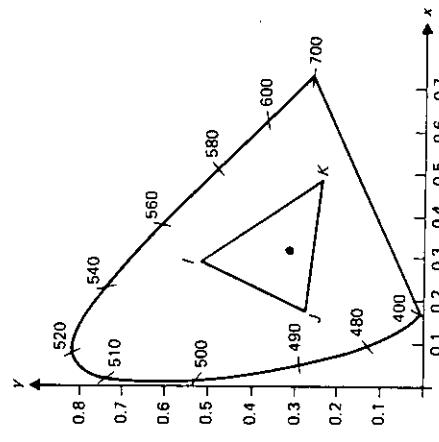


Fig. 17.25 Mixing three colors. All colors in triangle can be created.

Another use for the chromaticity diagram is to define and compare gamuts available from various color display and hardcopy devices. Color Plate 22 shows the gamuts for color TV, film, and print. The smallness of the print gamut with respect to the TV gamut suggests that if images originally created on a color TV must be faithfully reproduced by printing, a reduced gamut of colors should be used with the TV. Otherwise, accurate reproduction will not be possible. If, however, the goal is simply to make a pleasing rather than an exact reproduction, small differences in color gamuts are inconsequential.

17.4 COLOR MODELS FOR RASTER GRAPHICS

The purpose of a color model is to allow convenient specifications of colors within some color gamut. Our prime interest is the gamut for color TV, as defined by the RGB (red, green, blue) TV primaries in Color Plate 22. A secondary interest is the color gamut for hard-copy devices. A color model is a specification of a 3D color coordinate system and a 3D subspace in the coordinate system within which each playable color is represented by a point. The color models discussed here are all based on the RGB primaries, although in general any three primaries could be used. The models specify only colors in the RGB gamut.

Three hardware-oriented models are the RGB, used with color TV monitors; YIQ, which is the broadcast TV color system; and CMY (cyan, magenta, yellow) for color printing devices. Unfortunately, none of these models are particularly easy for a programmer or application user to control, because they do not directly relate to our intuitive color notions of hue, saturation, and brightness. Therefore, another class of models has been developed with ease of use as a goal. Several such models exist; we shall discuss only two: the HSV and HLS models. These and other models are described in [GSPC79, JOBL78, MEYE80, and SMIT78].

17.4.1. The RGB Color Model

The red, green, and blue color model uses a cartesian coordinate system. The subspace of interest is the unit cube, shown in Fig. 17.26. The RGB primaries are additive, i.e., the individual contributions of each primary are added together to form a result. The main diagonal of the cube, with equal amounts of each primary, represents the gray levels. Color Plate 23 shows several views of RGB color space. The RGB model is of interest primarily because it is used in color TV monitors and in many raster displays. There is also a considerable body of knowledge concerning the eye's response and sensitivity to colors specified as (R, G, B) triples.

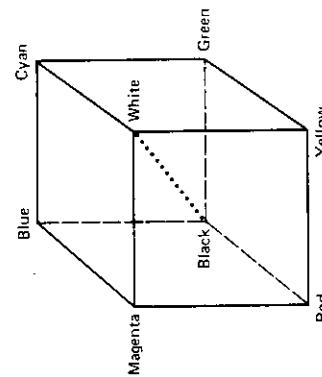


Fig. 17.26 RGB color cube. Grays are on the dotted main diagonal.

17.4.2 The CMY Color Model

Cyan, magenta, and yellow are the complements of red, green, and blue, respectively. They are called *subtractive primaries* because their effect is to subtract some color from white light. The subspace of the cartesian coordinate system for CMY is the same as for RGB, except that white (full light) is at the origin instead of black (no light). Colors are specified by what is removed or subtracted from white light, rather than by what is added to blackness.

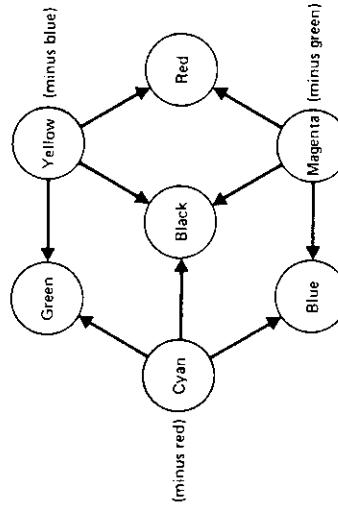
A knowledge of CMY is important when dealing with hard-copy devices which deposit colored pigments onto paper, such as the Xerox copier and the Applicon ink jet plotter. When a surface is coated with cyan ink or paint, no red light is reflected from the surface. Cyan subtracts red from the reflected white light, which is itself the sum of red, green, and blue. Hence, in terms of the additive primaries, cyan is blue plus green. Similarly, magenta absorbs green so it is red plus blue, while yellow absorbs blue so it is red plus green. A surface coated with cyan and yellow absorbs red and blue, leaving only green to be reflected from illuminating white light. A cyan, yellow, and magenta surface absorbs red, green, and blue, and therefore is black. These relations, diagrammed in Fig. 17.27, can be seen in Color Plate 23 and are represented by the equations:

$$\begin{bmatrix} C \\ M \\ Y \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} R \\ G \\ B \end{bmatrix}. \quad (17.13)$$

The unit column vector is the RGB representation for white and the CMY representation for black. Note that the Y in CMY is not the same as the Y in the CIE system.

The conversion from RGB to CMY is then

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} C \\ M \\ Y \end{bmatrix}. \quad (17.14)$$



17.4.3 The YIQ Color Model

The YIQ model is important because it is used in commercial color TV broadcasting and is therefore closely related to color raster graphics. YIQ is a recoding of RGB for transmission efficiency and for downward compatibility with black and white TV. The Y in YIQ is in fact the same as the Y in the CIE's X , Y , and Z primaries: a primary whose spectral energy distribution matches the luminosity response curve. The Y component of a color TV signal is shown on black and white TV. The YIQ model is in a 3D cartesian coordinate system, with the subspace being the convex polyhedron that maps into the RGB cube. The RGB to YIQ conversion which performs the mapping is defined as:

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.30 & 0.59 & 0.11 \\ 0.60 & -0.28 & -0.32 \\ 0.21 & -0.52 & 0.31 \end{bmatrix} \cdot \begin{bmatrix} R \\ G \\ B \end{bmatrix}. \quad (17.15)$$

The inverse of the matrix is used for the reverse conversion. Color Plate 24 shows planes of constant Y in the YIQ model. Further discussion of YIQ can be found in [SMIT78, PRIT77].

By specifying colors with the YIQ model, an important TV problem can be avoided: two colors that look quite different to our eyes may look exactly the same when transmitted or videotaped and then viewed on a black and white monitor. This can be avoided by guaranteeing that two colors that are meant to be distinguished one from the other (such as a filled area and its border) have different luminances (values of Y), so they are displayed at different intensities.

The YIQ model is designed to exploit a useful property of our visual system, which is more sensitive to changes in luminosity than to changes in hue or saturation. This suggests using more bits (or bandwidth) to represent Y than to represent I and Q , thus providing higher resolution in Y . Also, objects that cover a very small part of our field of view produce no color sensation, but are perceived only by their intensity. This suggests that less spatial resolution is needed for I and Q than for Y . The NTSC [PRIT77] encoding of YIQ into a broadcast signal uses both these properties to maximize the amount of information transmitted in a fixed bandwidth.

17.4.4 The HSV Color Model

The RGB, CMY, and YIQ models are hardware oriented. By contrast, Smith's HSV (hue, saturation, value) model [SMIT78] is user oriented, being based on the intuitive appeal of the artist's tint, shade, and tone. The subspace within which the model is defined is a hexcone, or six-sided cone, as in Fig. 17.28. The top of the hexcone corresponds to $V = 1$, which contains the maximum-value (intensity) colors. Note that complementary colors are 180° opposite one another as measured by H . This is the angle around the vertical axis, with red at 0° . The value of S is a ratio, ranging from 0 on the center line (V -axis) to 1 on the triangular sides of the hexcone. Saturation is measured relative to the gamut represented by the model, not relative to the CIE chart, and thus is not the same as purity.

Fig. 17.27 Subtractive primaries (cyan, magenta, yellow) and their mixtures.

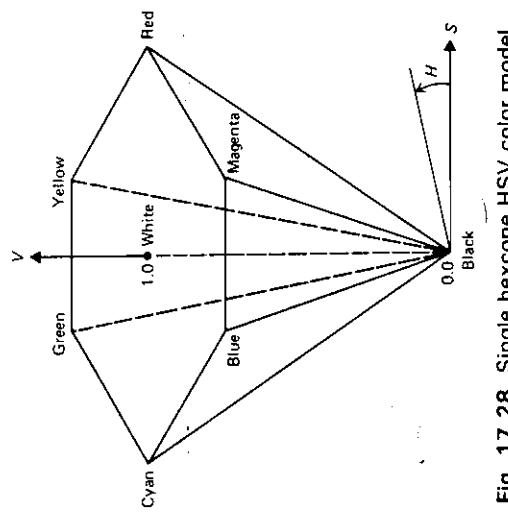


Fig. 17.28 Single hexcone HSV color model.

The hexcone is one unit high in V , with the apex at the origin. The point at the apex is black and has a coordinate of $V = 0$. Any value of S between 0 and 1 can be associated with the point $V = 0$. The point $S = 0$, $V = 1$ is white. Intermediate values of V for $S = 0$ (on the center-line) are the grays. When $S = 0$, the value of H is irrelevant and is called *undefined*. When S is not zero, H is relevant. For example, pure red is at $H = 0$, $S = 1$, $V = 1$. Indeed, any color with $V = 1$, $S = 1$ is akin to an artist's pure pigment used as the starting point in mixing colors. Adding white pigment corresponds to decreasing S (without changing V). Adding black pigment corresponds to decreasing V (without changing S). Tones are created by decreasing both S and V . Of course changing H corresponds to selecting the pure pigment with which to start. Thus H , S , and V each correspond one-to-one to concepts from the artists' color system.

The top of the HSV hexcone corresponds to the surface seen by looking along the principal diagonal of the RGB color cube from white toward black. Such a view is shown in Fig. 17.29 and in Color Plate 23(c). The RGB cube has subcubes, as illustrated in Fig. 17.30 on page 616. Each subcube, when viewed along its main diagonal, appears like the hexagon in Fig. 17.29, except smaller. Each plane of constant V in HSV space corresponds to such a view of a subcube in RGB space. The main diagonal of RGB space becomes the V axis of HSV space. Thus we see intuitively the correspondence between RGB and HSV. The following two algorithms define the correspondence precisely by providing conversions from one model to the other:

```

procedure RGB_TO_HSV(r, g, b: real; var h, s, v: real)
{Given: r, g, b, each in [0, 1]}
{Desired: h in [0, 360], s and v in [0, 1], except if s = 0,
then h = undefined which is a defined constant whose value is outside the
interval [0, 360]}
begin
  max := MAXIMUM(r, g, b);
  min := MINIMUM(r, g, b);
  v := max;
  if max <> 0
    then s := (max - min)/max
    else s := 0;
  if s = 0
    then h := undefined
  else
    begin
      rc := (max - r)/(max - min); {rc measures "distance" of color
from red}
      gc := (max - g)/(max - min);
      bc := (max - b)/(max - min);
      if r = max then h := bc - gc {resulting color between
yellow and magenta}
      else if g = max then h := 2 + rc - bc {resulting color between cyan
and yellow},
      else if b = max then h := 4 + gc - rc; {resulting color between
magenta and cyan}
      h := h*60;
      if h < 0 then h := h + 360
    end {chromatic case}
  end {RGB_TO_HSV}

```

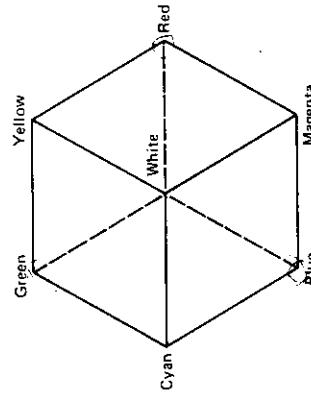


Fig. 17.29 RGB color cube viewed along principal diagonal. Visible edges of cube are solid, while the invisible ones are dashed.

```

procedure HSV_TO_RGB(var r, g, b: real; h, s, v: real);
{Given: h in [0, 360] or undefined, s and v in [0, 1]}
{Desired: r, g, b, each in [0, 1]}
begin
  if s = 0
    then
      {achromatic color: there is no hue}
      if h = undefined
        then
          {this is the achromatic case}
          begin
            r := v;
            g := v;
            b := v
          end
        else ERROR
      else
        begin
          if h = 360 then h = 0;
          h := h/60;
          i := FLOOR(h);
          f := h - i;
          p := v*(1 - s);
          q := v*(1 - (s*f));
          t := v*(1 - (s*(1 - f)));
          case i of
            0: (r, g, b) := (v, t, p);
            1: (r, g, b) := (q, v, p);
            2: (r, g, b) := (p, v, t);
            3: (r, g, b) := (p, q, v);
            4: (r, g, b) := (t, p, v);
            5: (r, g, b) := (v, p, q);
          end
        end
      end
    {hue}
end
{HSV_TO_RGB}

```

17.4.5 The HLS Color Model

The HLS (hue, lightness, saturation) color model, used by Tektronix and based on the Ostwald [OSTW31] color system, forms the double hexcone subspace seen in Fig. 17.31. Hue is the angle around the vertical axis of the double hexcone, with red at 0° .^{*} The colors occur around the perimeter in the same order they occur in the CIE diagram when its boundary is traversed counterclockwise: red, yellow, green, cyan, blue, and magenta. This is also the same order as in the HSV single hexcone model. In fact, one can think of HLS as a deformation of HSV, in which white is “pulled” upwards to form the upper hexcone from the $V = 1$ plane. As with the single hexcone model, the complement of any hue is located 180° further around the double hexcone, and saturation is measured radially from the vertical axis, from 0 on the axis to 1 on the surface. Lightness is 0 for black (at the lower tip of the double hexcone) to 1 for white (at the upper tip). Color Plate 25 shows an exploded view of the HLS model, in which the double hexcone has been topologically deformed into a double cone.

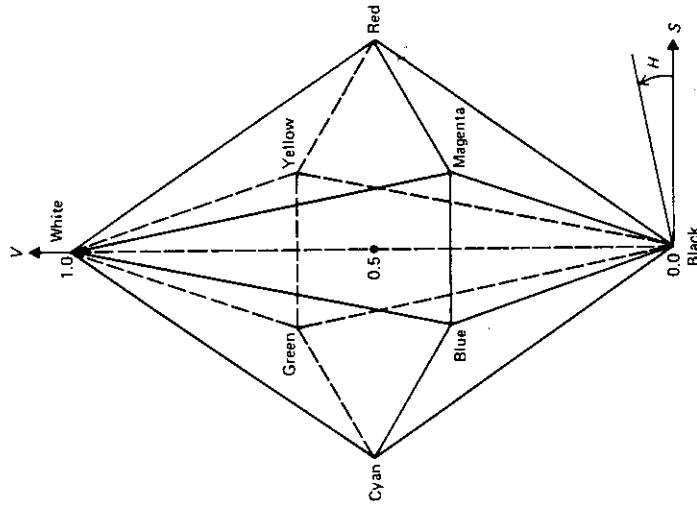


Fig. 17.31 Double hexcone HLS color model.

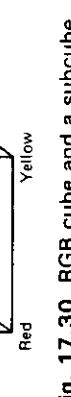


Fig. 17.30 RGB cube and a subcube.

^{*}For consistency with the HSV model, we have changed from the Tektronix convention of blue at 0° and depict the model as a double hexcone rather than as a double cone.

The HLS model, like the HSV model, is easy to use. The grays all have $S = 0$, but the maximally saturated hues are at $S = 1$, $L = 0.5$. If potentiometers are used to specify the color-model parameters, setting $L = 0.5$ to get the strongest possible colors is a disadvantage over the HSV model, in which $S = 1$ and $V = 1$ achieve the same effect. Notice, however, that neither the V in HSV nor the L in HLS corresponds to luminance in the YIQ model, so two different colors defined in either space can easily have the same luminance and be indistinguishable on black and white TV or videotape. Also note that, strictly speaking, "lightness" is a term normally used in discussing light reflection, and hence is used somewhat loosely here.

The following conversion procedures are modified from those given by Metrick [GSPC79] to leave H as *undefined* when $S = 0$ and to have $H = 0$ for red rather than for blue:

```

procedure RGB_TO_HLS(r, g, b: real; var h, l, s: real);
{Given: r, g, b, each in [0, 1]}
{Desired: h in [0, 360), l and s in [0, 1], except if s = 0,
then h = undefined}
begin
  max := MAXIMUM(r, g, b);
  min := MINIMUM(r, g, b);
  l := (max + min)/2;
  {Calculate saturation}
  if max = min
    then
      begin
        r := g = b; {achromatic case}
      end
    else
      begin
        if l <= 0.5 then s := (max - min)/(max + min)
        else s := (max - min)/(2 - max - min);
        {Calculate hue}
        rc := (max - r)/(max - min);
        gc := (max - g)/(max - min);
        bc := (max - b)/(max - min);
        if r = max then h := bc - gc
        else if g = max then h := 2 + rc - bc
        else if b = max then h := 4 + gc - rc;
        h := h*60;
        if h < 0.0 then h := h + 360
      end
    end
  {RGB_TO_HLS}

```

```

procedure HLS_TO_RGB(var r, g, b: real; h, l, s: real);
{given: h in [0, 360] or undefined, l and s in [0, 1]}
{desired: r, g, b, each in [0, 1]}
function VALUE(n1, n2, hue)
begin
  if hue > 360 then hue := hue - 360;
  if hue < 0 then hue := hue + 360;
  if hue < 60 then VALUE := nl + (n2 - nl)*hue/60;
  else if hue < 180 then VALUE := n2;
  else if hue < 240 then VALUE := nl + (n2 - nl)*(240 - hue)/60;
  else
    begin
      VALUE := nl
      if l <= 0.5 then m2 := l*(1 + s)
      else m2 := l + s - l*s;
      nl := 2*m1 - m2;
      if s = 0
        then
          begin
            if h = undefined
              then r := g := b := l
            else ERROR
          end
        begin
          r := VALUE(nl, m2, h + 120);
          g := VALUE(nl, m2, h);
          b := VALUE(nl, m2, h - 120)
        end
      end
    end
  {HLS_TO_RGB}

```

17.4.6 Interpolating in Color Space

Color interpolation is necessary in at least three cases: Gouraud shading (Section 16.4), blending two images together as in a "fade-in", "fade-out" sequence, and combining the color of a partially transparent surface with that of another surface (Section 16.6). The results of the interpolation depend on the color model in which the colors are interpolated; hence care must be taken to select an appropriate model. If the conversion from one color model to another transforms a straight line (representing the interpolation path) in one color model into a straight line in the other color model, the interpolation results in both models will be the same. This is the case for the RGB, CMY, and YIQ color models, all of which are related by simple affine transformations. However, a straight line in the RGB model does *not* in general transform into a straight line in either the HSV or HSL models. Consider, for example, the interpolation between red and green. In RGB, Red = (1, 0, 0) and Green = (0, 1, 0). Their interpolation (with both weights equal to 0.5 for convenience) is (0.5, 0.5, 0). Applying algorithm *RGB_TO_HSV* to this result, we have (60°, 1, 0.5). Now, representing red and green in HSV, we have (0°, 1, 1) and

$(120^\circ, 1, 1)$. But interpolating with equal weights in HSV we have $(60^\circ, 1, 1)$; thus the value differs by 0.5 from the same interpolation in RGB.

As a second example, consider interpolating red and cyan in both the RGB and HSV models. In RGB, we start with $(1, 0, 0)$ and $(0, 1, 1)$, respectively, and interpolate to $(0.5, 0.5, 0.5)$ which in HSV is represented as *(undefined, 0, 0.5)*. In HSV, red and cyan are $(0^\circ, 1, 1)$ and $(180^\circ, 1, 1)$. Interpolating, we have $(90^\circ, 1, 1)$: a new hue at maximum value and saturation has been introduced, whereas the “right” result from combining equal amounts of complementary colors is a gray value. Here again we have a case where interpolating and then transforming gives different results than transforming and then interpolating.

Which model should be used for interpolation? If the traditional results from additive colors are desired (note that interpolation is basically an additive process), then RGB is preferred to HSV or HLS. If, on the other hand, the objective is to interpolate between two colors of fixed hue (or saturation) and to maintain the fixed hue (saturation) for all interpolated colors, then HSV or HLS is preferable.

17.5 REPRODUCING COLOR HARD COPY

Color images are reproduced in print in a way similar to that used for monochrome images, but four sets of halftone dots are printed, one for each of the subtractive primaries, plus black. Black is used because in printing it is difficult to obtain a deep black by combining the three primaries. The dots are carefully positioned with respect to one another so as not to overlap. The orientation of each of the grids of dots is different, to avoid creating interference patterns. Color Plate 26 shows an enlarged halftone color pattern. Our eyes spatially integrate the light reflected from adjacent dots, so we see the color defined by the proportions of primaries in adjacent dots. This spatial integration of different colors is the same phenomenon we experience when viewing the triads of red, green, and blue dots on color TV.

We deduce, then, that color reproduction in print and on CRTs depends on the same spatial integration used in monochrome reproduction. The monochrome dithering techniques discussed in Section 17.2.2 can also be used with color to extend the number of available colors, again at the expense of resolution. Consider a color display with three bits per pixel—one each for red, green, and blue. We can use a 2×2 pixel pattern area to obtain 125 different colors, because each pattern can display five intensities for each of red, green, and blue, by using the halftone patterns in Fig. 17.8. This results in $5 \times 5 \times 5 = 125$ color combinations. Color Plate 27 shows 125 colors generated in exactly this way.

Not all color reproduction depends exclusively on spatial integration. For instance, the Xerox color copier and Appleton ink jet plotter actually mix subtractive pigments on the paper’s surface to obtain a small set of different colors. In the case of the Xerox copier, the colored pigments are first deposited in three successive steps, then heated and melted together. The inks sprayed by the plotter mix before drying. Spatial integration may be used to further expand the color range.

17.6 COLOR HARMONY

Contemporary color display and hard-copy devices can produce a wide gamut of colors. Some color pairs harmonize well with one another, while other pairs clash with great dissonance. How can we select colors that harmonize? Many books have been written on color selection, including [BIRR61]; we state here a few of the simpler rules which will help produce color harmony.

The most fundamental rule is to use colors selected according to some method, typically by traversing a smooth path in a color model and/or by restricting the colors to planes (or hexcones) of constant value in a color model. This might mean using colors of the same hue, or two complementary colors and mixtures thereof, or colors of constant lightness or value. For instance, Color Plate 28 shows the use of colors chosen along a path in the RGB color model. Furthermore, colors are best spaced at equal *perceptual* distances in whatever subspace they are drawn from (this is not the same as being equally spaced increments of a coordinate in the subspace and can be difficult to implement). Note too that linear interpolation (as in Gouraud shading) between two colors produces different results in different color spaces (see Exercise 14).

A random selection of different hues and saturations will usually appear quite garish. Alvy Ray Smith described an experiment in which a 16×16 grid was filled with randomly generated colors. Not unexpectedly, the grid was unattractive. Sorting the 256 colors according to their H , S , and V and redisplaying them on the grid in their new order gave a remarkable improvement to the appearance of the grid. More specific instances of these rules suggest that if a chart contains just a few colors, the complement of one of the colors should be used in the background. With an image containing many different colors, a neutral (gray) background should be used. If two adjoining colors are not particularly harmonious, a thin black border can be used to set them apart.

From a physiological viewpoint, the eye’s low sensitivity to blue suggests that blue on a black background will be hard to distinguish. By inference, yellow (the complement of blue) on a white background (the complement of black) also will be relatively hard to distinguish (see Exercise 10). For the sake of those of us who are red-green color-blind (the most common form), it is also good to avoid reds and greens with low saturation and luminance.

17.7 USING COLOR IN INTERACTIVE GRAPHICS

There are two groups of people who must deal with color in interactive graphics systems: the programmer and the user. The programmer’s job is to provide the user with an understandable color model or reasonable selection of colors. In simple applications, the programmer will assign colors when creating an image—the user will not be involved. This might be the case in command and control or data presentation applications. In other applications, the user will simply be presented with a

fixed set of colors from which to select, typically in the form of a grid of color squares displayed across the bottom of the screen. A logical pick device is employed to select the desired color. Sometimes several sets of colors (usually called palettes) might be available, each one being a set of colors which have been chosen, by the rules of color harmony, as an attractive color scheme.

The most general application (with respect to color) allows the user to create a palette of colors, either by specification of points in some color space or constructively by blending a few base colors in various ways to create new hues and to make tints, tones, and shades of the hues. This would typically be done in painting and animation applications, where very fine control over color is necessary.

EXERCISES

- 17.1 Derive an equation for the number of intensities that can be represented by $m \times m$ pixel patterns, where each pixel has w bits.
- 17.2 Write the programs needed to gamma-correct a black and white display through a look-up table. Input parameters are γ , I_0 , m , the number of intensities desired, and c , the constant in Eq. (17.3).
- 17.3 Write a general algorithm to display a pixel array on a bi-level output device. The inputs to the algorithm are: (a) an $m \times m$ array of pixel intensities, with w bits per pixel; and (b) an $n \times n$ growth sequence matrix. Assume that the output device has resolution of $m \cdot n \times m \cdot n$.
- 17.4 Repeat the previous problem by using ordered dither. Now the output device has resolution $m \times m$, the same as the input array of pixel intensities.
- 17.5 Write an algorithm to display a filled polygon on a bi-level device by using an $n \times n$ filtering pattern.
- 17.6 If certain patterns are used to fill a polygon being displayed on an interlaced raster display, all of the "on bits" will fall either on the odd or the even scan lines, introducing a slight amount of flicker. Revise the algorithm from Exercise 5 to permute rows of the $n \times n$ pattern so that alternate replications of the pattern will alternate use of the odd and even scan lines. Figure 17.32 shows the results obtained by using intensity level one from Fig. 17.8, with and without this alternation.

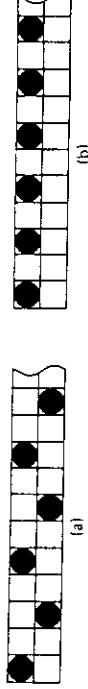


Fig. 17.32 Results obtained by using intensity level 1 from Fig. 17.8 in the following way: (a) with alternation (intensified pixels are on both scan lines), and (b) without alternation (all intensified pixels are on one scan line).